

On the Complexity of Closest Pair via Polar-Pair of Point-Sets

Roe David* Karthik C. S.† Bundit Laekhanukit‡

Weizmann Institute of Science§

Abstract

Every graph G can be represented by a collection of equi-radii spheres in a d -dimensional metric Δ such that there is an edge uv in G if and only if the spheres corresponding to u and v intersect. The smallest integer d such that G can be represented by a collection of spheres (all of the same radius) in Δ is called the *sphericity* of G , and if the collection of spheres are non-overlapping, then the value d is called the *contact-dimension* of G . In this paper, we study the geometric representation of a complete bipartite graph $K_{n,n}$ in various L^p metrics and connect it to the complexity of the closest pair problem. We prove the following results.

1. We show that contact-dimension and sphericity of $K_{n,n}$ in the L^0 -metric is n , and that the contact-dimension and sphericity of $K_{n,n}$ in L^p -metrics with $p > 2$ is $\Theta(\log n)$.
2. We give alternate proofs for the lower bound on the sphericity and the upper bound on the contact-dimension of $K_{n,n}$ in the L^2 -metric using spectral analysis. Our upper and lower bounds are slightly weaker than that in [Frankl & Maehara, DCG'88] and [Maehara, DCG'91], respectively. However, our proofs are arguably simpler than previous works including [Bilu & Linial, JCTB'05], which also uses spectral analysis.

All our upper bounds are constructive, and gives efficient algorithms to construct a *polar-pair* of point-sets, which are point-sets corresponding to the bipartition of $K_{n,n}$. As a result, we are able to show quadratic-time hardness for the closest pair problem in the L^p -metric with $d = \omega(\log n)$ and $p > 2$ under the Orthogonal Vectors Hypothesis. The hardness result is obtained by a reduction from the bichromatic variant of the closest pair problem in the Hamming metric whose hardness was recently established in [Alman & Williams, FOCS'15]. In addition, we prove subquadratic-time hardness for the closest pair (and thus bichromatic closest pair) problem in the L^∞ metric, which does not follow from [Alman & Williams, FOCS'15].

*This work was partially supported by the ISF (grant No. 621/12) and by the I-CORE Program (grant No. 4/11).

†This work was partially supported by Irit Dinur's ERC-StG grant number 239985.

‡This work was partially supported by the ISF (grant No. 621/12) and by the I-CORE Program (grant No. 4/11).

§roee.david@weizmann.ac.il, karthik.srikanta@weizmann.ac.il, bundit.laekhanukit@weizmann.ac.il

1 Introduction

This paper studies the geometric representation of a bipartite complete graph in L^p -metrics and connects it to the complexity of the closest pair problem. Given a point-set P in a d -dimensional L^p -metric, an α -distance graph is a graph $G = (V, E)$ with a vertex set $V = P$ and an edge set

$$E = \{ab : \|a - b\|_p \leq \alpha; a, b \in P; a \neq b\}.$$

In other words, points in P are centers of spheres of radius α , and G has an edge ab if and only if the spheres centered at a and b intersect. The *sphericity* of a graph G in an L^p -metric, denoted by $\text{sph}_p(G)$, is the smallest dimension d such that G is isomorphic to some α -distance graph in a d -dimensional L^p -metric, for some constant $\alpha > 0$. The sphericity of a graph in the L^∞ -metric is known as *cubity*. A notion closely related to sphericity is *contact-dimension*, which is defined in the same manner except that the spheres representing G must be non-overlapping. To be precise, an α -contact graph $G = (V, E)$ of a point-set P is an α -distance graph of P such that every edge ab of G has the same distance (i.e., $\|a - b\|_p = \alpha$). Thus, G has the vertex set $V = P$ and has an edge set E such that

$$\forall ab \in E, \quad \|a - b\|_p = \alpha \quad \text{and} \quad \forall ab \notin E, \quad \|a - b\|_p > \alpha.$$

The contact-dimension of a graph G in the L^p -metric, denoted by $\text{cd}_p(G)$, is the smallest dimension such that G is isomorphic to a contact-graph in the d -dimensional L^p -metric. We will use distance and contact graphs to means 1-distance and 1-contact graphs.

We are interested in determining the sphericity and the contact-dimension of the biclique $K_{n,n}$ in various L^p -metrics. For notational convenience, we denote $\text{sph}_p(K_{n,n})$ by $\text{bsph}(\Delta)$, the *biclique sphericity* of a metric Δ , and denote $\text{cd}_p(K_{n,n})$ by $\text{bcd}(\Delta)$, the *biclique contact-dimension* of Δ . We call a pair of point-sets (A, B) *polar* if it is the partition of a contact graph isomorphic to $K_{n,n}$. More precisely, a pair of point-sets (A, B) is polar in an L^p -metric if there exists a constant $\alpha > 0$ such that every inner-pair $a, a' \in A$ (resp., $b, b' \in B$) has L^p -distance greater than α while every crossing-pair $a \in A, b \in B$ has L^p -distance exactly α .

The biclique sphericity and contact-dimension of the L^2 and L^∞ metrics are well-studied in literature (see [Rob69, Mae84, Mae85, FM88, Mae91, BL05]). Maehara [Mae91, Mae84] showed that $n < \text{bsph}(L^2) \leq (1.5)n$, and Maehara and Frankl & Maehara [Mae85, FM88] showed that $(1.286)n - 1 < \text{bcd}(L^2) < (1.5)n$. The lower bound proofs in [Mae91] and [FM88] are long and involved. Bilu and Linial [BL05] gave an alternate proof for the linear lower bound on the biclique sphericity of the L^2 -metric using spectral analysis of matrices. The proof in [BL05] is shorter but gives a slightly weaker bound. For cubity, Roberts [Rob69] showed that $\text{bcd}(L^\infty) = \text{bsph}(L^\infty) = 2 \log_2 n$. Our main results extend these studies to other L^p metrics and also reflects upon known bounds (see Theorem 1).

Next, we connect the biclique contact-dimension to the complexity of the *closest pair* problem (CLOSEST PAIR), which asks to find a closest pair of points from a given collection P of m points in a d -dimensional metric space. Specifically, we tie the complexity of the CLOSEST PAIR to its bichromatic variant, namely, the *bichromatic closest pair* problem (BCP), in which a point in the collection P is colored by either red or blue, and the goal is to find a closest pair of red-blue points. Alman and Williams [AW15] showed that BCP admits no subquadratic-time algorithm for high

dimensions under the Orthogonal Vectors Hypothesis. Hence, it is a good candidate to base the complexity of CLOSEST PAIR on.

It is not hard to see that BCP is at least as hard as CLOSEST PAIR since we can apply an algorithm for BCP to solve CLOSEST PAIR with the same asymptotic running time. However, it is not clear whether the other direction is true. We will give a simple reduction from BCP to CLOSEST PAIR using a polar-pair of point-sets. First, take a polar-pair (A, B) , each with cardinality $n = m/2$, in the L^p -metric. Next, pair up vectors in A and B to red and blue points, respectively, and then attach a vertex $a \in A$ (resp., $b \in B$) to its matching red (resp., blue) point. This reduction increases the distances between every pair of points, but by the definition of the polar-pair, this process has more affect on the distances of the *monochromatic* (i.e., red-red or blue-blue) pairs than that of *bichromatic pairs*, and the reduction, in fact, has no effect on the order of crossing-pair distances at all. Consequently, by scaling the vectors in A and B appropriately, the reduction gives an instance of CLOSEST PAIR whose closest pair of points must be a bichromatic pair. Thus, we have that BCP and CLOSEST PAIR in the L^p metric are computationally equivalent for dimensions $d = \Omega(\text{bcd}(L^p))$.

Nevertheless, the hardness of BCP based on the Orthogonal Vectors Hypothesis need not be carried over to CLOSEST PAIR as the polar-pair (A, B) might have large dimension. Thus, we need an *efficient* construction of a polar-pair with $n^{o(1)}$ -dimension to show that BCP and CLOSEST PAIR in the considered L^p -metric are equivalent in the quadratic-time complexity.

1.1 Our Results and Contributions

Our main results are lower and upper bounds on the biclique-sphericity for the following metric spaces: L^0, L^1, L^2 and L^p with $p > 2$.

Theorem 1. *The following are upper and lower bounds of biclique-sphericity for the L^p metric.*

$$\text{bsph}(L^0) = \text{bcd}(L^0) = n \tag{1}$$

$$n \leq \text{bsph}(L_{\{0,1\}}^0) \leq \text{bcd}(L_{\{0,1\}}^0) \leq n^2 \quad (i.e., P \subseteq \{0,1\}^d) \tag{2}$$

$$\Omega(\log n) \leq \text{bsph}(L^1) \leq \text{bcd}(L^1) \leq O(n^2) \tag{3}$$

$$(n-3)/2 \leq \text{bsph}(L^2) \leq \text{bcd}(L^2) \leq 2n-1 \tag{4}$$

$$\text{bsph}(L^p) = \Theta(\text{bcd}(L^p)) = \Theta(\log n) \quad \text{for } p > 2 \tag{5}$$

Note that $\text{bsph}(\Delta) \leq \text{bcd}(\Delta)$ for any metric Δ . Thus, it suffices to prove a lower bound for $\text{bsph}(\Delta)$ and prove an upper bound for $\text{bcd}(\Delta)$. In (1), (2), (3), and (5), we give bounds on sphericity and contact dimension for the respective metrics. For the L^2 -metric, we give alternate proofs for linear lower and upper bounds for $\text{bsph}(L^2)$ and $\text{bcd}(L^2)$ respectively. While our lower bounds are slightly weaker than the best known bounds [FM88, Mae91], our proofs are much shorter (the original proofs take more than ten pages while our proofs only take a couple of pages). Our proofs use spectral analysis similar to that in [BL05], but most arguments require no heavy machinery and thus are arguably simpler than the previous works [FM88, Mae91, BL05]. Additionally, while our upper bound in (4) is weaker than the one in [Mae85], our construction yields a more structured embedding: every pair of points not corresponding to an edge are also equidistant.

Next, we establish the following hardness result for CLOSEST PAIR under the Orthogonal Vectors Hypothesis (OVH) [Wil05] (a weaker assumption than the Strong Exponential Time Hypothesis [IP01, IPZ01, CIP09]).

Corollary 2. *For the L^p metric with $p > 2$ and the L^∞ metric, there is a linear-time reduction from the bichromatic closest pair problem in the d -dimensional L^p -metric (resp., L^∞ -metric) to the closest pair problem in the $(d + O(\log n))$ -dimension L^p -metric (resp., L^∞ -metric).*

Recently, Alman and Williams [AW15] showed that, for any $\varepsilon > 0$, $p \geq 0$ and $d = \omega(\log n)$, BCP in the d -dimensional L^p -metric admits no $(n^{2-\varepsilon})$ -time algorithm unless the OVH is false. Thus, Corollary 2 implies the subquadratic-time hardness of CLOSEST PAIR in L^p -metric with $p > 2$ in our case.

Corollary 3. *For any $\varepsilon > 0$, $p > 2$ and $d = \omega(\log n)$, the closest pair problem in the d -dimensional L^p -metric admits no $(n^{2-\varepsilon})$ -time algorithm unless the Orthogonal Vectors Hypothesis is false.*

We remark that the hardness of BCP (and thus CLOSEST PAIR) in the L^∞ -metric does not follow from [AW15]. Nonetheless, we prove the quadratic-time hardness of CLOSEST PAIR directly in Theorem 4.

Theorem 4. *For any $\varepsilon > 0$ and $d = \omega(\log n)$, the closest pair problem in the d -dimensional L^∞ -metric admits no $(n^{2-\varepsilon})$ -time algorithm unless the Orthogonal Vectors Hypothesis is false.*

Additionally, for independent interest, we prove the quadratic-time hardness for a generalization of CLOSEST PAIR, namely, the *set closest pair* problem (SET CLOSEST PAIR) in which we are given a collection of point-sets (instead of points), and the goal is to find a closest pair of point-sets (two sets that minimizes the set-distances). Note that the distance between two sets U and V is the minimum distance over all pairs $u \in U$ and $v \in V$, i.e., the distance of the closest pair between the sets. This computation seems to be inherent in many problems of machine learning, and in similarity searches in particular. Thus, the hardness result for SET CLOSEST PAIR might be of interest to the machine learning community. To be precise, we prove the following hardness result:

Theorem 5. *For any $\varepsilon > 0$, $p > 1$, and $d = \omega(\log n)$, the set closest pair problem in the d -dimensional L^p -metric admits no $(n^{2-\varepsilon})$ -time algorithm unless the Orthogonal Vectors Hypothesis is false.*

We remark that unlike CLOSEST PAIR, we could prove conditional hardness result for SET CLOSEST PAIR in the *Euclidean metric*.

1.2 Related Works

While our paper studies sphericity and contact-dimension of the complete bipartite graph, determining the contact-dimension of a complete graph in L^p -metrics has also been extensively studied in the notion of *equilateral dimension*. To be precise, the equilateral dimension of a metric Δ which is the maximum number of equidistant points that can be packed in Δ . An interesting connection is in the case of the L^1 -metric, for which we are unable to establish a strong lower bound for $\text{bsph}(L^1)$. The equilateral dimension of L^1 is known to be at least $2d$, and this bound is believed

to be tight [Guy83]. This is a notorious open problem known as *Kusner's conjecture*, which is confirmed for $d = 2, 3, 4$ [BCL98, KLS00], and the best upper bound for $d \geq 5$ is $O(d \log d)$ by Alon and Pavel [AP03]. If Kusner's conjecture is true for all d , then $\text{sph}_1(K_n) = n/2$.

The complexity of CLOSEST PAIR has been a subject of study for many decades. There have been a series of developments on CLOSEST PAIR in the Euclidean space (see, e.g., [Ben80, HNS88, KM95, SH75, BS76]), which culminates in a deterministic $O(2^{O(d)} n \log n)$ -time algorithm [BS76] and a randomized $O(2^{O(d)} n)$ -time algorithm [Rab76, KM95]. For low (i.e., constant) dimensions, these algorithms are tight as the matching lower bound of $\Omega(n \log n)$ was shown by Ben-Or [Ben83] and Yao [Yao91] for the *algebraic decision tree* model, thus settling the complexity of CLOSEST PAIR in low dimensions. For high dimension (i.e., $d = \omega(\log n)$), there is a huge gap between the upper and lower bounds. There is no known algorithm that runs in time significantly better than a trivial $O(n^2 d)$ -time algorithm for general d except for the case that $d \geq \Omega(n)$, wherein there are subcubic algorithms in L^1 and L^∞ metrics [GS16, ILLP04]. Nevertheless, the best known lower bound for CLOSEST PAIR was $\Omega(\max\{n \log n, nd\})$ [Ben83, Yao91] (note that nd is the input size). There were no better lower bounds for any L^p -metrics prior to our work.

1.3 Organization

This paper is organized as follows. In Section 3, we prove lower and upper bounds on $\text{bsph}(L^2)$ and $\text{bcd}(L^2)$ (Eq. 4 in Theorem 1). In Section 4, we discuss the case of the L^1 -metric. Although we are unable to prove either a strong lower or upper bound for this case, we do make progress towards proving lower bound for $\text{bsph}(L^1)$ and $\text{bcd}(L^1)$. In Section 5, we prove *tight* bounds for $\text{bsph}(L^0)$ and $\text{bcd}(L^0)$ (Eq. 1 in Theorem 1). In Section 6, we prove logarithmic bounds on $\text{bsph}(L^p)$ and $\text{bcd}(L^p)$, for $p > 2$ (Eq. 5 in Theorem 1). In Section 7, we use the algorithmic construction of upper bounds in Section 6 to prove subquadratic hardness for CLOSEST PAIR and SET CLOSEST PAIR. Finally, in Section 8, we conclude our paper by giving open problems and highlighting some directions for future research.

2 Preliminaries

We use the following standard terminologies and notations. Below we define below L^p -norms, which are distance measures in the metric spaces considered in this paper.

Distance Measures. For any vector $x \in \mathbb{R}^d$, the L^p -norm of x is denoted by $\|x\|_p = \left(\sum_{i=1}^d |x_i|^p\right)^{1/p}$. The L^∞ -norm of x is denoted by $\|x\|_\infty = \max_{i \in [d]} \{|x_i|\}$, and the L^0 -norm of x is denoted by $\|x\|_0 = |\{x_i \neq 0 : i \in [d]\}|$, i.e., the number of non-zero coordinates of x . These norms define distance measures of \mathbb{R}^d . The distance of two points x and y w.r.t. the L^p -norm is, thus, $\|x - y\|_p$. The distances measures that are well studied in literature are the *Hamming distance* L^0 -norm, the *Rectilinear distance* L^1 -norm, the *Euclidean distance* L^2 -norm, the *Chebyshev distance* (a.k.a, *Maximum-norm*) L^∞ -norm. Given two sets of points $A, B \subseteq \mathbb{R}^d$, the L^p -distance between the two sets is defined to be $d_p(A, B) = \min_{a \in A, b \in B} \|a - b\|_p$.

Problems. We define CLOSEST PAIR as a decision problem in a usual manner. In particular, in CLOSEST PAIR, we are given a real number α and a collection of points $P \subseteq \Sigma^d$ in a d -dimensional L^p -metric, for some finite set Σ , and the goal is to determine if there exists a pair of distinct points $a, b \in P$ such that

$$\|a - b\|_p \leq \alpha.$$

For BCP, we are given a real α , and we specify the two color classes (the collections of red and blue points) by A and B , and the goal is to determine if there exists a pair of points $a \in A$ and $b \in B$ such that

$$\|a - b\|_p \leq \alpha.$$

A generalization of CLOSEST PAIR that we study is the *set closest pair* problem (SET CLOSEST PAIR). In SET CLOSEST PAIR, we are given a real α and a collection (possible, multiset) \mathcal{X} of N subsets of points in a d -dimensional metric space, and the goal is to determine if there exists a pair of distinct sets $A, B \in \mathcal{X}$ such that

$$\text{dist}_p(A, B) \leq \alpha,$$

where $\text{dist}_p(A, B) = \min_{a \in A, b \in B} \|a - b\|_p$.

Note that the search variant of the above problems are at least as hard as the decision variants described above, and thus any lower bounds on the above problems also applies to their search version.

Fine-Grained Complexity and Conditional Hardness. Conditional hardness is the current trend in proving running-time lower bounds for polynomial-time solvable problems. This has now developed into the area of *Fine-Grained Complexity*. Please see, e.g., [Wil15, Wil16] and references therein.

The *Orthogonal Vectors Hypothesis* (OVH) is a popular complexity theoretic assumption in fine-grained complexity to base the hardness of a problem on. OVH states that in the Word RAM model with $O(\log n)$ bit words, any algorithm requires $n^{2-o(1)}$ time in expectation to determine whether collections of vectors $U, W \subseteq \{0, 1\}^d$, each with $n/2$ vectors, where $d = \omega(\log n)$, contains an orthogonal pair $u \in U$ and $w \in W$ (i.e., $\sum_{i=1}^d u_i \cdot w_i = 0$).

Another popular conjecture is the *Strong Exponential-Time Hypothesis* for SAT (SETH), which states that, for any $\varepsilon > 0$, there exists an integer k , such that k -SAT on n variables cannot be solved in $O(2^{(1-\varepsilon)n})$ -time. Williams showed that SETH implies OVH [Wil05].

3 Geometric Representation of Biclique in L^2

In this section, we prove a lower bound on $\text{bsph}(L^2)$ and an upper bound on $\text{bcd}(L^2)$.

3.1 Lower Bound on the Biclique-Sphericity

We show using spectral analysis that $\text{bsph}(L^2) \geq (n - 3)/2$.

Theorem 6. *For every $n, d \in \mathbb{N}$, and any two sets $A, B \subseteq \mathbb{R}^d$, each of cardinality n , suppose the following holds for some non-negative real numbers a and b with $a > b$.*

1. For every u and v both in A we have that $\|u - v\|_2 > a$.
2. For every u and v both in B we have that $\|u - v\|_2 > a$.
3. For every u in A and v in B , we have that $\|u - v\|_2 \leq b$.

Then the dimension d must be at least $\frac{n-3}{2}$.

Proof. Let $|A| = |B| = n$ be arbitrary two sets of vectors in \mathbb{R}^d that satisfy the above conditions. We will show that $d \geq \frac{n-3}{2}$. First, we scale all the vectors in $A \cup B$ so that the vector with the largest ℓ_2 -norm in $A \cup B$ has ℓ_2 -norm that is equal to 1 (by this scaling, the parameters a, b are scaled as well by, say, s . For brevity, we will write a for a/s and similarly for b). We modify A and B in two steps as follows. We add one new coordinate to all of the vectors with value $K \gg 1$ (to be determined exactly later) and obtain $A_1, B_1 \subseteq \mathbb{R}^{d+1}$. Note that each element in the new set of vectors A_1 and B_1 has ℓ_2 -norm roughly equal to K . More specifically, the square of the ℓ_2 -norm is bounded between K^2 and $K^2 + 1$ and the vector with the largest ℓ_2 -norm in $A_1 \cup B_1$ has ℓ_2 -norm that is equal to $\sqrt{K^2 + 1}$.

By adding to the last coordinate of each vector u in $A_1 \cup B_1$ a positive value c_u smaller than $1/K$, we can impose that all the vectors are with ℓ_2 -norm that is equal to $\sqrt{K^2 + 1}$. To see this, note that if we have a vector u_1 in $A_1 \cup B_1$ that has ℓ_2 -norm that is equal to K (namely, as small as possible), then by setting c_{u_1} to satisfy

$$(K + c_{u_1})^2 = K^2 + 1, \quad (6)$$

we get that the ℓ_2 -norm of u_1 is $\sqrt{K^2 + 1}$. Any c_{u_1} that solves Equation 6 is smaller than $1/K$ and by assuming that u_1 has a larger ℓ_2 -norm we would have get better bounds on c_{u_1} .

Let $A'_1 \cup B'_1$ be the set of vectors that was obtained by adding c_u 's as described above. Let u, v be vectors in $A_1 \cup B_1$ and let u', v' be the corresponding vectors in $A'_1 \cup B'_1$. By definition, the following holds:

$$\begin{aligned} \|u - v\|_2^2 &\leq \|u' - v'\|_2^2 = \|u - v\|_2^2 + (c_u - c_v)^2 \\ &\leq \|u - v\|_2^2 + 1/K^2. \end{aligned}$$

Hence, by choosing K to satisfy $1/K^2 \leq \frac{a^2 - b^2}{2}$ it follows that $A'_1 \cup B'_1$ satisfies the conditions of the theorem with $a' = a$ and $b' = \sqrt{b^2 + \frac{a^2 - b^2}{2}} < a'$. Again, for brevity, we refer to a' as a and b' as b .

Given $A'_1, B'_1 \subseteq \mathbb{R}^{d+1}$, let a_1, a_2, \dots, a_n be the vectors from A'_1 , and b_1, b_2, \dots, b_n be the vectors from B'_1 . Consider the following matrix in $\mathbb{R}^{2(d+1) \times 2n}$:

$$M = \begin{pmatrix} a_1, a_2, \dots, a_n & b_1, b_2, \dots, b_n \\ b_1, b_2, \dots, b_n & a_1, a_2, \dots, a_n \end{pmatrix} \quad (7)$$

Define the set A_2 to be the first n column vectors of M and define B_2 to be the last n column vectors of M . Note that $A_2 \cup B_2 \subseteq \mathbb{R}^{2(d+1)}$ and satisfies the conditions of the theorem with

$a_2 = 2a_1 > 2b_1 = b_2$. Consider the inner product matrix $M^T M \in \mathbb{R}^{2n \times 2n}$ written in a block matrix form as follows:

$$M^T M = cI_{2n \times 2n} + \begin{pmatrix} M_{1,1} & M_{1,2} \\ M_{2,1} & M_{2,2} \end{pmatrix},$$

where $M_{1,1}, M_{1,2}, M_{2,1}, M_{2,2} \in \mathbb{R}^{n \times n}$ and c is such that the matrix $\begin{pmatrix} M_{1,1} & M_{1,2} \\ M_{2,1} & M_{2,2} \end{pmatrix}$ has the value 0 on the diagonal elements (recall that all the vectors have the same ℓ_2 -norm). By the definition of M (see Equation 7), one can check that the following hold.

1. The matrices $M_{1,1}, M_{1,2}, M_{2,1}, M_{2,2}$ are all symmetric: for $M_{1,1}, M_{2,2}$ it follows since $M^T M$ is a symmetric matrix, and for $M_{1,2}, M_{2,1}$ it follows by the way M was defined; see Equation 7.
2. $M_{1,1} = M_{2,2}$. This follows by Equation 7.
3. $M_{1,2} = M_{2,1}$. This follows since $M_{1,2} = M_{2,1}^T = M_{2,1}$. Here the first equality follows since $M^T M$ is a symmetric matrix, and the last equality follows by item 1.

Hence we can write $M^T M = cI_{2n \times 2n} + \begin{pmatrix} M_{1,1} & M_{1,2} \\ M_{1,2} & M_{1,1} \end{pmatrix}$. In the rest of the proof, we analyze some of the eigenvectors of $M^T M$, for this we consider the matrix $M_{1,1} - M_{1,2}$. Since both $M_{1,1}$ and $M_{1,2}$ are symmetric, we have that $M_{1,1} - M_{1,2}$ is symmetric as well and has real eigenvalues. Moreover, by the conditions of the theorem, it holds that $M_{1,1} - M_{1,2}$ is strictly negative (i.e., all the entries of the matrix are negative). This follows because all the vectors have the same ℓ_2 -norm. Let x_1, x_2, \dots, x_n be the eigenvectors of $M_{1,1} - M_{1,2}$ with eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$. By the Perron–Frobenius Theorem it follows that λ_1 is strictly smaller than $\lambda_2, \lambda_3, \dots, \lambda_n$.

Let $x_i \in \mathbb{R}^n$ be an eigenvector of $M_{1,1} - M_{1,2}$ with eigenvalue λ_i . Then the following holds.

$$\begin{aligned} \begin{pmatrix} M_{1,1} & M_{1,2} \\ M_{1,2} & M_{1,1} \end{pmatrix} \begin{pmatrix} x_i \\ -x_i \end{pmatrix} &= \begin{pmatrix} (M_{1,1} - M_{1,2}) x_i \\ -(M_{1,1} - M_{1,2}) x_i \end{pmatrix} \\ &= \begin{pmatrix} \lambda_i x_i \\ -\lambda_i x_i \end{pmatrix} \\ &= \lambda_i \begin{pmatrix} x_i \\ -x_i \end{pmatrix}. \end{aligned}$$

Hence the vectors $\begin{pmatrix} x_1 \\ -x_1 \end{pmatrix}, \begin{pmatrix} x_2 \\ -x_2 \end{pmatrix}, \dots, \begin{pmatrix} x_n \\ -x_n \end{pmatrix}$ are eigenvectors of $\begin{pmatrix} M_{1,1} & M_{1,2} \\ M_{1,2} & M_{1,1} \end{pmatrix}$ with eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$. The operation of adding $cI_{2n \times 2n}$ to $\begin{pmatrix} M_{1,1} & M_{1,2} \\ M_{1,2} & M_{1,1} \end{pmatrix}$ shifts the eigenvalues of $M^T M$ to $\lambda_1 + c, \lambda_2 + c, \dots, \lambda_n + c$.

Note that $M^T M$ is a positive semidefinite matrix. Hence, $\lambda_1 + c, \lambda_2 + c, \dots, \lambda_n + c \geq 0$. More specifically, $\lambda_1 + c \geq 0$ and $\lambda_2 + c, \dots, \lambda_n + c > 0$ (since $\lambda_1 < \lambda_2, \lambda_3, \dots, \lambda_n$). It follows that $M^T M$ has at least $n - 1$ positive eigenvalues. Hence, the rank of $M^T M$ is at least $n - 1$. By standard linear algebra arguments, it holds that the rank of M is at least the rank of $M^T M$, and the rank of M is at most $2(d + 1)$. That is,

$$2(d + 1) \geq \text{rank}(M) \geq \text{rank}(M^T M) \geq n - 1.$$

Therefore, $d \geq (n - 3)/2$, and the theorem follows. \square

3.2 Upper Bound on the Biclique Contact-Dimension

Now we show that $\text{bcd}(L^2) \leq 2n - 1$. Our proof is constructive and yields a linear-time algorithm that constructs a polar-pair (A, B) with *strong* property: the distance of every inner pair $a, a' \in A$ (resp., $b, b' \in B$) is the same.

Theorem 7. *For every $n \in \mathbb{N}$, there exists two sets $A, B \subseteq \mathbb{R}^{2n-1}$ each of cardinality n such that the following holds:*

1. *For every u and v both in A , we have $\|u - v\|_2 = \sqrt{2}$.*
2. *For every u and v both in B , we have $\|u - v\|_2 = \sqrt{2}$.*
3. *For every u in A and v in B , we have $\|u - v\|_2 = \sqrt{2} \cdot (1 - 1/n)^{1/2}$.*

Moreover, there exists a deterministic algorithm that outputs A and B and runs in linear-time w.r.t. the size of the output (which is $\Theta(n^2)$).

Proof. Let $G = K_{n,n}$ be the complete bipartite graph. Consider the adjacency matrix A_G of the graph G and note that $n \cdot I + A_G$ is a positive semi-definite matrix (this follows since the smallest eigenvalue of G is $-n$). Let U be the symmetric matrix $n \cdot I + A_G$ and W be a matrix with $2n$ columns that satisfies $W^T W = U$. One can check that the rank of W in this case is $2n - 1$. We construct A and B as follows. The first n columns of W are the vectors of A and the last n columns of W are the vectors of B . If both $u, v \in A$ (and similarly for B) then we note the following:

$$\begin{aligned} \|v - u\|_2^2 &= (U)_{v,v} - 2(U)_{u,v} + (U)_{u,u} \\ &= 2n. \end{aligned}$$

On the other hand, if $u \in A$ and $v \in B$, then we have the following:

$$\begin{aligned} \|v - u\|_2^2 &= (U)_{v,v} - 2(U)_{u,v} + (U)_{u,u} \\ &= 2n - 2. \end{aligned}$$

The following claim shows that we can construct W , and consequently, A and B efficiently (note that, in general, factoring a positive semi-definite matrix takes $O(n^3)$).

Claim 8. *The matrix W can be constructed in $O(n^2)$ time.*

Proof. Let $\vec{1}$ be the all one column vector and $\vec{1}_P$ be a column vector with values 1 on the coordinates of A and -1 on the coordinates of B . Note that A_G is a rank 2 matrix because,

$$A_G = \frac{1}{2} \left(\vec{1} (\vec{1})^T - \vec{1}_P (\vec{1}_P)^T \right).$$

One can check that $\vec{1}$ and $\vec{1}_P$ are eigenvectors of A_G with eigenvalues n and $-n$, respectively. The other eigenvectors are orthogonal to $\vec{1}$ and $\vec{1}_P$, and their associated eigenvalues are 0 (because A_G has rank 2).

Recall that the Hadamard vectors H_1, H_2, \dots, H_n of dimension n is a collection of n vectors in $\{1, -1\}^n$ with the property that every two vectors are orthogonal and H_1 is the all one vector. There exist well known recursive constructions of the Hadamard vectors (that are linear in the output size) when n is a power of 2. In the case that it is not a power of 2, it suffices to work with the smallest power of 2 which is greater than n .

Consider the following matrix:

$$V_{2n \times 2n} = \begin{pmatrix} \vec{1} & H_2 & H_3 & \dots & H_n & \vec{0} & \vec{0} & \dots & \vec{0} & \vec{1}_P \\ & \vec{0} & \vec{0} & \dots & \vec{0} & H_2 & H_3 & \dots & H_n & \end{pmatrix}.$$

We note that the column vectors $\vec{v}_1, \vec{v}_2, \dots, \vec{v}_{2n}$ of V are the eigenvectors of A_G . Furthermore, U has the same eigenvectors (with n added to the eigenvalues). So, if $\vec{v}_1, \vec{v}_2, \dots, \vec{v}_{2n}$ are the normalized eigenvectors of U with positive eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_{2n-1}$ and $\lambda_{2n} = 0$, then we have

$$W^T = [\vec{v}_1, \vec{v}_2, \dots, \vec{v}_{2n}] \text{diag} \left\{ \sqrt{\lambda_1}, \sqrt{\lambda_2}, \dots, \sqrt{\lambda_{2n-1}} \right\}.$$

The last matrix multiplication can be done in $O(n^2)$ since $\text{diag} \left\{ \sqrt{\lambda_1}, \sqrt{\lambda_2}, \dots, \sqrt{\lambda_{2n}} \right\}$ is a diagonal matrix. \square

The proof follows by appropriately scaling the vectors in A and B . \square

4 The Possibility of Representing Biclique in L^1

In this section, we discuss the case of the L^1 -metric, which is the only case that we are unable to prove neither strong lower bound nor linear upper bound. A weak lower bound $\text{bsph}(L^1) \geq \Omega(\log n)$ follows from the proof for the L^p -metric with $p > 2$ in Section 6.1, and a quadratic upper bound $\text{bcd}(L^1) \leq n^2$ follows from the proof for the L^0 -metric in Section 6.2. However, we cannot prove any upper bound smaller than $\Omega(n^2)$ or any lower bound larger than of $O(\log n)$. There is a huge gap in this case. Hence, we study an average case relaxation of the question.

We show in Theorem 10 that there is no distribution whose expected distances simulate a polar-pair of point-sets in the L^1 -metric. Consequently, even though we could not prove the biclique sphericity lower bound for the L^1 -metric, we are able to refute an existence of an geometric representation with *large gap* for any dimension as shown in Corollary 11. (A similar result was shown in [DM94].)

Definition 9 (L^1 -distribution). *Let X, Y be two random variables taking values in \mathbb{R}^d . An L^1 -distribution is constructed by X, Y if the followings hold.*

$$\mathbb{E}_{x_1, x_2 \in_R X} [\|x_1 - x_2\|_1] > \mathbb{E}_{x_1 \in_R X, y_1 \in_R Y} [\|x_1 - y_1\|_1], \quad (8)$$

$$\mathbb{E}_{y_1, y_2 \in_R Y} [\|y_1 - y_2\|_1] > \mathbb{E}_{x_1 \in_R X, y_1 \in_R Y} [\|x_1 - y_1\|_1]. \quad (9)$$

Lemma 10. *For any two finite-supported random variables X, Y that are taking values in \mathbb{R}^d there is no L^1 -distribution.*

Proof. Assume towards a contradiction that there exist two finite-supported random variables X, Y that are taking values in \mathbb{R}^d and that are satisfying Inequality 8 of Definition 9. Given a vector $x \in \mathbb{R}^d$, we denote by $x(i)$ the value of the i -th coordinate of x . Hence the following inequalities hold,

$$\begin{aligned}
0 &> \mathbb{E}_{x_1 \in_R X, y_1 \in_R Y} [\|x_1 - y_1\|_1] - \mathbb{E}_{x_1, x_2 \in_R X} [\|x_1 - x_2\|_1] \\
&= \mathbb{E}_{x_1, x_2 \in_R X, y_1 \in_R Y} [\|x_1 - y_1\|_1 - \|x_1 - x_2\|_1] \\
&= \frac{1}{d} \mathbb{E}_{x_1, x_2 \in_R X, y_1 \in_R Y} \left[\mathbb{E}_{i \in [1 \dots d]} [|x_1(i) - y_1(i)| - |x_1(i) - x_2(i)|] \right] \\
&= \frac{1}{d} \mathbb{E}_{i \in [1 \dots d]} \left[\mathbb{E}_{x_1, x_2 \in_R X, y_1 \in_R Y} [|x_1(i) - y_1(i)| - |x_1(i) - x_2(i)|] \right].
\end{aligned}$$

Thus for some $i^* \in [d]$ the following holds,

$$0 > \mathbb{E}_{x_1, x_2 \in_R X, y_1 \in_R Y} [|x_1(i^*) - y_1(i^*)| - |x_1(i^*) - x_2(i^*)|]. \quad (10)$$

Fix $i^* \in [d]$ satisfying the above inequality. For the sake of clarity, we assume that the random variables X, Y are taking values in \mathbb{R} (i.e., projection on the i^{th} coordinate). We can assume that the size of $\text{supp}(X) \cup \text{supp}(Y)$ is greater than 1 because, if $\text{supp}(X) \cup \text{supp}(Y)$ contains a single point then $\mathbb{E}_{x_1 \in_R X, y_1 \in_R Y} [\|x_1 - y_1\|_1] = \mathbb{E}_{x_1, x_2 \in_R X} [\|x_1 - x_2\|_1] = 0$, contradicting Inequality 10. Let $\text{supp}(X) \cup \text{supp}(Y)$ contains $t \geq 2$ points. We prove by induction on t , that there are no X, Y over \mathbb{R} satisfying Inequality 10. The induction base case is when $t = 2$. By Inequality 10, there exists 3 points $\tilde{x}_1, \tilde{x}_2, \tilde{y}_1$ in \mathbb{R} such that,

$$0 > \|\tilde{x}_1 - \tilde{y}_1\|_1 - \|\tilde{x}_1 - \tilde{x}_2\|_1. \quad (11)$$

Since $\text{supp}(X) \cup \text{supp}(Y)$ contains exactly two points then $\tilde{x}_1, \tilde{x}_2, \tilde{y}_1$ are supported by two distinct points in \mathbb{R} . Hence, there are two cases, either that $x_1 = x_2$ (and $y_1 \neq x_1$), or that $x_1 \neq x_2$ (and either $\tilde{y}_1 = \tilde{x}_1$ or $\tilde{y}_1 = \tilde{x}_2$). It is easy to see that none of these cases satisfy Inequality 11, resulting in a contradiction.

We assume from induction hypothesis that there are no X, Y taking values over \mathbb{R} satisfying Inequality 10 when the size of $\text{supp}(X) \cup \text{supp}(Y)$ is equal to $k \geq 2$. We consider the case when $t = k + 1 \geq 3$. Sort the points in $\text{supp}(X) \cup \text{supp}(Y)$ by their value and denote by s_i the value of the i -th point of $\text{supp}(X) \cup \text{supp}(Y)$. For the sake of simplicity we say that we *change* the value of s_{t-1} to \tilde{s}_{t-1} , where $s_{t-2} \leq \tilde{s}_{t-1} \leq s_t$, if after changing its value we change the values of (at least one of) X, Y to \tilde{X}, \tilde{Y} in such a way that the value of the $t - 1$ -th point (after sorting) of $\text{supp}(\tilde{X}) \cup \text{supp}(\tilde{Y})$ is equal to \tilde{s}_{t-1} (if $s_{t-2} = \tilde{s}_{t-1}$ then the value of the $(t - 2)$ -th point of $\text{supp}(\tilde{X}) \cup \text{supp}(\tilde{Y})$ is equal to \tilde{s}_{t-1}). Define the function $f : [s_{t-2}, s_t] \rightarrow \mathbb{R}$ as follows:

$$f(x) = \mathbb{E}_{x_1 \in_R \tilde{X}, y_1 \in_R \tilde{Y}} [\|x_1 - y_1\|_1] - \mathbb{E}_{x_1, x_2 \in_R \tilde{X}} [\|x_1 - x_2\|_1],$$

where \tilde{X}, \tilde{Y} are obtained after changing s_{t-1} to $x \in [s_{t-2}, s_t]$. The crucial observation is that the function f is linear. Hence either $f(s_{t-2}) \geq f(s_{t-1})$ or $f(s_t) \geq f(s_{t-1})$ and we can reduce the size of $\text{supp}(X) \cup \text{supp}(Y)$ by 1. However, this contradicts our induction hypothesis. \square

The next lemma refutes the existence of a polar-pair of point-sets with large gap in any dimension.

Corollary 11 (No Polar-Pair of Point-Sets in L^1 with Large Gap). *For any $\alpha > 0$, there exist no subsets $A, B \subseteq \mathbb{R}^d$ of $n/2$ vectors with $d < n/2$ such that*

- For any u, v both in A , or both in B , we have $\|u - v\|_1 \geq \frac{1}{1-2/n} \cdot \alpha$.
- For any $u \in A$ and $v \in B$, $\|u - v\|_1 < \alpha$.

Proof. Assume towards a contradiction that there exist a polar-pair of point-sets (A, B) in the L^1 -metric that satisfies the conditions above. We can create a distribution X and Y such that

$$\mathbb{E}_{x_1, x_2 \in_R X} [\|x_1 - x_2\|_1] = \mathbb{E}_{y_1, y_2 \in_R Y} [\|y_1 - y_2\|_1] > \mathbb{E}_{x \in_R X, y \in_R Y} [\|x - y\|_1]$$

To see this, we create a uniform random variable X (resp., Y) for the set A (resp., B). Now the expected distance of two independent copies of X (resp., Y) is at least $\frac{1}{1-2/n} \alpha \cdot \left(1 - \frac{1}{n/2}\right) = \alpha$, which follows because we may pick the same point twice. Since the expected distance of the crossing pair $a \in A$ and $b \in B$ is less than α . This contradicts Theorem 10. \square

In stark contrast, a result similar to Theorem 10 but in the L^0 metric, i.e., Lemma 15 implies that there is no polar-pair in L^0 . Finally, one can show that biclique sphericity lower bound for L^2 implies that there is no polar-pair distribution for L^2 as well. This completes the picture of non-existence of polar-pair distributions for all the three metrics (i.e., L^2, L^1 , and L^0).

5 Geometric Representation of Biclique in L^0

In this section, we prove a lower bound on $\text{bsph}(L^0)$ and upper bounds on $\text{bcd}(L^0)$. The proofs in this section are mostly combinatorial due to the behavior of the Hamming metric.

We start by providing real-to-binary reductions in Section 5. Then we proceed to prove the lower bound on $\text{bsph}(L^0)$ in Section 5.1 and then the upper bounds on $\text{bcd}(L^0)$ in Section 5.2.

Real to Binary Reductions First we prove the following lemma, which allows mapping from vectors in \mathbb{R}^d to zero-one vectors.

Lemma 12 ((Trivial) Real to Binary Reduction). *Let $S \subseteq \mathbb{R}$ be a finite set of real numbers, and let $P \subseteq S^d$. Then there exists a transformation $\phi : S^d \rightarrow \{0, 1\}^{d|S|}$ such that, for any $x, y \in P$,*

$$\|x - y\|_0 = \|\phi(x) - \phi(y)\|_0$$

Proof. First, we order element in S in arbitrary order and write it as $S = \{r_1, r_2, \dots, r_{|S|}\}$. Next we define $\psi : S \rightarrow \{0, 1\}^{|S|}$ so that the i -th coordinate of $\psi(r_i)$ is 1 and others are zeroes. That is, $b = \psi(r_i)$ is such that

$$b_j = \begin{cases} 1 & \text{if } j = i \\ 0 & \text{otherwise} \end{cases}$$

Then we define $\phi(x) = (\psi(x_1), \psi(x_2), \dots, \psi(x_d))$. Clearly, $\|\psi(r_i) - \psi(r_j)\|_0 = 1$ if and only if $r_i \neq r_j$. Therefore, we conclude that for any $x, y \in P$,

$$\|\phi(x) - \phi(y)\|_0 = \|x - y\|_0.$$

□

Next we present a randomized reduction which gives a better dimension bound for some parameters. We require the additive form of Chernoff's bound.

Lemma 13 (Additive Chernoff's bound). *Let X_1, X_2, \dots, X_N be independent binary random variables such $\Pr[X_i = 1] = \Pr[X_i = 0] = 1/2$ for all $i \in [N]$. Then*

$$\Pr \left[\sum_i X_i > N/2 + a \right] < e^{a^2/(n/2)} \quad \text{and} \quad \Pr \left[\sum_i X_i < N/2 - a \right] < e^{a^2/(n/2)}.$$

Lemma 14 (Randomized Real to Binary Reduction with Given Bar). *Let $\beta > 0$ be any integer. Let $P \subseteq \mathbb{R}^d$ be a set of d -dimensional points with cardinality n . Then there exists a function $\phi : P \rightarrow \{0, 1\}^{dL}$, where $L > 8\beta \ln n$, that transforms points in P to binary strings such that, for any points $u, v, x, y \in P$,*

$$\|u - v\|_0 > \beta \geq \|x - y\|_0 \implies \|\phi(u) - \phi(v)\|_0 > \|\phi(x) - \phi(y)\|_0.$$

Proof. We prove the claim by probabilistic method.

First, observe that, for each coordinate, there can be at most n different values of p_i , for any vector $p \in P$. Thus, we may think that the i -th coordinate of any point in P is a letter from an alphabet set of size n , denoted by Σ_i . Then we construct a function $\phi_i : \Sigma_i \rightarrow \{0, 1\}^L$ for each alphabet set Σ_i by uniformly and independently at random setting each coordinate of $\phi_i(a)$, for $a \in \Sigma_i$, to be 0 and 1 with probability $1/2$. That is, each letter $a \in \Sigma_i$ is mapped to a binary string $b = \phi_i(a) \in \{0, 1\}^L$ in such a way that

$$b_j = \begin{cases} 1 & \text{with probability } 1/2 \\ 0 & \text{otherwise} \end{cases}$$

Note that we need to construct a mapping for each coordinate independently. The mapping ϕ is then defined as $\phi(x) = (\phi_1(x_1), \phi_2(x_2), \dots, \phi_d(x_d))$.

Let us take any points $u, v, x, y \in P$ such that $\alpha = \|u - v\|_0 > \beta \geq \|x - y\|_0$. (Thus, $\alpha \geq \beta + 1$). Observe that α (resp., β) is the number of coordinates that the vectors u and v (resp., x and y) are different.

Now consider $\|\phi(u) - \phi(v)\|_0$. By construction, for each coordinate i such that $u_i = v_i$, we must have $\phi_i(u) = \phi_i(v)$. Otherwise, if $\phi_i(u) \neq \phi_i(v)$, then each bit of the blocks $\phi_i(u)$ and $\phi_i(v)$ (i.e., $\phi_i(u)_j$ and $\phi_i(v)_j$) are different with probability $1/2$, and these events are independent. Thus,

$$\mathbb{E}[\|\phi(u) - \phi(v)\|_0] = \alpha \cdot L/2.$$

By Chernoff's bound (Lemma 13), we have

$$\Pr \left[\|\phi(u) - \phi(v)\|_0 < \alpha \cdot L/2 - (\alpha \cdot L \cdot 4 \ln n)^{1/2} \right] < \frac{1}{n^4}.$$

Next consider $\|\phi(x) - \phi(y)\|_0$. By the same arguments as that of u and v , we conclude that

$$\Pr \left[\|\phi(x) - \phi(y)\|_0 > \beta \cdot L/2 + (\beta \cdot L \cdot 4 \ln n)^{1/2} \right] < \frac{1}{n^4}.$$

Since $\alpha - \beta \geq 1$ and $L > 4\beta \ln n$, $\|\phi(u) - \phi(v)\|_0 > \|\phi(x) - \phi(y)\|_0$ with probability $1 - 2/n^4$.

Therefore, by taking union bound over all pairs, we conclude that there exists ϕ such that $\|\phi(u) - \phi(v)\|_0 > \|\phi(x) - \phi(y)\|_0$ for all pairs u, v, x, y with $\|u - v\|_0 > \|x - y\|_0$. \square

5.1 Lower Bound on the Biclique-Sphericity

Now we will show that $\text{bsph}(L^0) \geq n$. Our proof requires the following lemma, which rules out a randomized algorithm that generates a polar-pair of point-sets.

Lemma 15 (No Distribution for L^0). *For any $\alpha > \beta \geq 0$, regardless of dimension, there exist no distributions \mathcal{A} and \mathcal{B} of points in \mathbb{R}^d with finite supports such that*

- $\mathbb{E}_{x, x' \sim \mathcal{A}}[\|x - x'\|_0] \geq \alpha$.
- $\mathbb{E}_{y, y' \sim \mathcal{B}}[\|y - y'\|_0] \geq \alpha$.
- $\mathbb{E}_{x \sim \mathcal{A}, y \sim \mathcal{B}}[\|x - y\|_0] \leq \beta$.

Proof. We prove by contradiction. Assume to a contrary that such distributions exist. Then

$$\mathbb{E}_{x, x' \sim \mathcal{A}}[\|x - x'\|_0] + \mathbb{E}_{y, y' \sim \mathcal{B}}[\|y - y'\|_0] - 2\mathbb{E}_{x \sim \mathcal{A}, y \sim \mathcal{B}}[\|x - y\|_0] > 0. \quad (12)$$

Let A and B be supports of \mathcal{A} and \mathcal{B} , respectively. Since \mathcal{A} and \mathcal{B} have finite supports, we may assume wlog that A and B are multi-sets of equal size. Moreover, by Lemma 12, we may further assume that vectors in A and B are binary vectors.

Observe that each coordinate of vectors in A and B contribute to the expectations independently. In particular, Eq. (12) can be written as

$$2 \sum_i \rho_{0,i}^A \rho_{1,i}^A + 2 \sum_i \rho_{0,i}^B \rho_{1,i}^B + 2 \sum_i (\rho_{0,i}^A \rho_{1,i}^B + \rho_{0,i}^B \rho_{1,i}^A) > 0 \quad (13)$$

where $\rho_{0,i}^A$, $\rho_{1,i}^A$, $\rho_{0,i}^B$ and $\rho_{1,i}^B$ are the probability that the i -th coordinate of $x \in A$ (resp., $y \in B$) is 0 (resp., 1). Thus, to show a contradiction, it is sufficient to consider the coordinate which

contributes the most to the summation in Eq. (13). The contribution of this coordinate to the summation is

$$2\rho_0^A \rho_1^A + 2\rho_0^B \rho_1^B - 2(\rho_0^A \rho_1^B + \rho_1^A \rho_0^B) = 2(\rho_0^A(\rho_1^A - \rho_1^B) + \rho_1^B(\rho_0^B - \rho_0^A)) = 2(\rho_0^A - \rho_0^B)(\rho_1^A - \rho_1^B) \quad (14)$$

Since $\rho_0^A + \rho_1^A = 1$ and $\rho_0^B + \rho_1^B = 1$, the summation in Eq.(14) can be non-negative only if $\rho_0^A = \rho_0^B$ and $\rho_1^A = \rho_1^B$. But, then this implies that the summation in Eq.(14) is zero. We have a contradiction since this coordinate contributes the most to the summation in Eq. (13) which we assume to be positive. \square

The next Lemma shows that $\text{bsph}(L^0) \geq n$.

Theorem 16 (Lower Bound for L^0 with Arbitrary Alphabet). *For any integers $\alpha > \beta \geq 0$ and $n > 0$, there exist no subsets $A, B \subseteq \mathbb{R}^d$ of n vectors with $d < n$ such that*

- For any $a, a' \in A$, $\|a - a'\|_0 \geq \alpha$.
- For any $b, b' \in B$, $\|a - a'\|_0 \geq \alpha$.
- For any $a \in A$ and $b \in B$, $\|a - b\|_0 < \beta$.

Proof. Suppose for a contradiction that such subsets A and B exist with $d < n$. We build uniform distributions \mathcal{A} and \mathcal{B} by uniformly at random picking a vector in A and B , respectively. Then it is easy to see that the expected value of inner distance is

$$\mathbb{E}_{x, x' \sim \mathcal{A}}[\|x - x'\|_0] \geq \alpha - \frac{\alpha}{n}$$

The intra distance of B is similar. We know that $\alpha - \beta \geq 1$ because they are integers and so are ℓ_0 -distances. But, then if $\alpha < n$, we would have distributions that contradict Lemma 15. Note that α and β are at most d (dimension). Therefore, we conclude that $d \geq n$. \square

5.2 Upper Bound on the Biclique Contact-Dimension

Now we show that, for vectors in \mathbb{R}^d , $\text{bcd}(L^0) \leq n$.

Theorem 17 (Upper Bound for L^0 with Arbitrary Alphabet). *For any integer $n > 0$ and $d = n$, there exist subsets $A, B \subseteq \mathbb{R}^d$ each with n vectors such that*

- For any $a, a' \in A$, $\|a - a'\|_0 = d$.
- For any $b, b' \in B$, $\|a - a'\|_0 = d$.
- For any $a \in A$ and $b \in B$, $\|a - b\|_0 = d - 1$.

Proof. First we construct a set of vectors A . For $i = 1, 2, \dots, n$, we define the i -th vector a of A so that a is an all- i vector. That is,

$$a = (i, i, \dots, i)$$

Next we construct a set of vectors B . The first vector of B is $(1, 2, \dots, n)$. Then the $i + 1$ -th vector of B is the left rotation of the i -th vector. Thus, the i -th vector of B is

$$b = (i, i + 1, \dots, n, 1, 2, \dots, i - 1).$$

It can be seen that the L^0 distance between any two vectors from the same set is d because all the coordinates are different. Any vectors from different set, say $a \in A$ and $b \in B$, must have at least one common coordinate. Thus, their L^0 distance is $d - 1$. This proves the lemma. \square

Below is the upper bounds for zero-one vectors, which is a corollary of Theorem 17.

Corollary 18 (Upper Bound for L^0 with Binary Vectors). *For any integer $n > 0$ and $d = n^2$, there exist subsets $A, B \subseteq \mathbb{R}^d$ each with n vectors such that*

- For any $a, a' \in A$, $\|a - a'\|_0 = n$.
- For any $b, b' \in B$, $\|a - a'\|_0 = n$.
- For any $a \in A$ and $b \in B$, $\|a - b\|_0 = n - 1$.

Proof. We take the construction from Theorem 17 and denote the two sets by A' and B' and denote their dimension by $d' = n$.

We transform A' and B' to sets A and B by applying the transformation ϕ in Lemma 12. That is,

$$A = \{\phi(a) : a \in A'\} \quad \text{and} \quad B = \{\phi(b) : b \in B'\}.$$

Since the alphabet set in Lemma 12 is $[n]$, we have a construction of A and B with dimension $d = n^2$. \square

The corollary below follows directly from Theorem 17 and Lemma 12.

Corollary 19 (Upper Bound for ℓ_0 with Binary Vectors). *For any integer $n > 0$ and $d = n^2$, there exist subsets $A, B \subseteq \mathbb{R}^d$ each with n vectors such that*

- For any $a, a' \in A$, $\|a - a'\|_0 = n$.
- For any $b, b' \in B$, $\|a - a'\|_0 = n$.
- For any $a \in A$ and $b \in B$, $\|a - b\|_0 = n - 1$.

It is not hard to see that the following corollary holds.

Corollary 20 (Upper Bound for L^0 with Small Alphabet Set). *For any integers $n > 0$, $K > 0$ and $d = n^2/K$, there exist subsets $A, B \subseteq ([K] \cup \{0\})^d$ each with n vectors such that*

- For any $a, a' \in A$, $\|a - a'\|_0 = d$.
- For any $b, b' \in B$, $\|a - a'\|_0 = d$.

- For any $a \in A$ and $b \in B$, $\|a - b\|_0 = d - 1$.

Proof Sketch. First, we follow the proof of Lemma 12. Here we can reduce the dimension by a factor of K because each coordinate can represent K values. The corollary then follows by following the proof of Corollary 18. \square

6 Geometric Representation of Biclique in L_p for $p > 2$

In this section, we show the lower bound on $\text{bsph}(L^p)$ and an upper bound on $\text{bcd}(L^p)$ for $p > 2$. Both bounds are logarithmic. The latter upper bound is constructive and thus implies the subquadratic-time equivalent between CLOSEST PAIR and BCP.

6.1 Lower Bound on the Biclique Sphericity

Now we show the lower bound on the biclique sphericity of a complete bipartite graph in L^p -metrics with $p > 2$. In fact, we prove the lower bound for the case of a star graph on n vertices, denoted by S_n , and then use the fact that $\text{bsph}(H) \leq \text{bsph}(G)$ for all induced subgraph H of G (i.e., $\text{bsph}(K_{n/2, n/2}, L^p) \geq \text{bsph}(S_{n/2}, L^p)$).

In short, we show in Lemma 23 that $O(\log n)$ is the maximum number of L_p -balls of radius $1/2$ that we can pack in an L_p -ball of radius 1 so that no two of them intersect or touch each other. This upper bounds, in turn, implies the lower bound on the dimension. We proceed the proof by volume arguments, which are commonly used in proving the minimum number of points in an ϵ -net that are sufficient to cover all the points in a sphere.

Definition 21 (ϵ -net). *The unit L_p -ball in \mathbb{R}^d centered at o is denoted by*

$$B(L_p^d, o) = \{x \in \mathbb{R}^d \mid \|x - o\|_p \leq 1\}.$$

For brevity, we write $B(L_p^d, o)$ as $B(L_p^d)$. Let (X, d) be a metric space and let S be a subset of X and ϵ be a constant greater than 0. A subset N_ϵ of X is called an ϵ -net of S under d if for every point $x \in S$ it holds for some point $y \in N_\epsilon$ that $d(x, y) \leq \epsilon$.

The following lemma is well known in literature, and we provide below a proof using essentially the same arguments as in [Ver10] for the sake of completeness.

Lemma 22. *There exists an ϵ -net for $B(L_p^d)$ under the L_p -metric of cardinality $(1 + \frac{2}{\epsilon})^d$.*

Proof. Let us fix $\epsilon > 0$ and choose N_ϵ of maximal cardinality such that $\|x - y\|_p > \epsilon$ for all $x \neq y$ both in N_ϵ . We claim that N_ϵ is an ϵ -net of the $B(L_p^d)$. Otherwise, there would exist an $x \in B(L_p^d)$ that is at least ϵ -far from all points in N_ϵ . Thus, $N_\epsilon \cup \{x\}$ contradicts the maximality of N_ϵ . After establishing that N_ϵ is an ϵ -net, we note that by the triangle inequality, we have that the balls of radii $\epsilon/2$ centered at the points in N_ϵ are disjoint. On the other hand, by the triangle inequality all such balls lie in $(1 + \epsilon/2) B(L_p^d)$. Comparing the volumes gives us that $\text{vol}(\frac{\epsilon}{2} B(L_p^d)) \cdot |N_\epsilon| \leq \text{vol}((1 + \frac{\epsilon}{2}) B(L_p^d))$. Since for all $r \geq 0$, we have that $\text{vol}(r \cdot B(L_p^d)) = r^d \cdot \text{vol}(B(L_p^d))$, we conclude the following: $|N_\epsilon| \leq \frac{(1 + \epsilon/2)^d}{(\epsilon/2)^d} = (1 + \frac{2}{\epsilon})^d$. \square

Theorem 23. *For every $N, d \in \mathbb{N}$, and any two sets $A, B \subseteq \mathbb{R}^d$, each of cardinality N , suppose the following holds for some non-negative real numbers a and b with $a > b$.*

1. *For every u and v both in A we have that $\|u - v\|_p > a$.*
2. *For every u and v both in B we have that $\|u - v\|_p > a$.*
3. *For every u in A and v in B , we have that $\|u - v\|_p \leq b$.*

Then the dimension d must be at least $\log_5(N)$.

Proof. Scale and translate the sets A, B in such a way that $b = 1$ and that $\vec{0} \in B$. It follows that $A \subseteq B(L_p^d)$. By Lemma 22 we can fix an $1/2$ -net $N_{1/2}$ for $B(L_p^d)$ of size 5^d . Note that for every $x \in N_{1/2}$ it holds that $1/2 \cdot B(L_p^d, x)$ contains at most one point from A . Note also that $N_{1/2}$ covers $B(L_p^d)$. Thus $|A| \leq 5^d$ which implies that $d \geq \log_5(N)$. \square

6.2 Upper Bound on the Biclique Contact-Dimension

We first give a simple randomized construction that gives a logarithmic upper bound on the biclique contact-dimension of L^p . The construction is simple. We uniformly at random take a subset A of n vectors from $\{-1, 1\}^{d/2} \times \{0\}^{d/2}$ and a subset B of n vectors from $\{0\}^{d/2} \times \{-1, 1\}^{d/2}$. Observe that, for any $p > 2$, the L_p -distance of any pair of vectors $a \in A$ and $b \in B$ is exactly d while the *expected distance* between the inner pair $a, a' \in A$ (resp., $b, b' \in B$) is strictly larger than d . Thus, if we choose d to be sufficiently large, e.g., $d \geq 10 \ln n$, then we can show by using a standard concentration bound (e.g., Chernoff's bound) that the probability that the inner-pair distance is strictly larger than d is at least $1 - 1/n^3$. Applying the union bound over all inner-pairs, we have that the d -neighborhood graph of $A \cup B$ is a bipartite complete graph with high probability. Moreover, the distances between any crossing pairs $a \in A$ and $b \in B$ are the same for all pairs. This shows the upper bound for the contact-dimension of a biclique in the L^p -metric for $p > 2$.

The above gives a simple proof of the upper bound on the biclique contact-dimension of the L^p -metric. Moreover, it shows a randomized construction of the polar-pair in $O(\log n)$ -dimensional L^p -metric, for $p > 2$, thus implying that CLOSEST PAIR and BCP are equivalent for these L^p -metrics.

For a deterministic construction, the proof is more involved, and we show it using appropriate codes.

Theorem 24. *Let $\zeta = 2/(2^{p-2} - 1)$. There exists two sets $|P| = |Q| = 2^{n/2}$ of vectors in $R^{2\alpha n}$, for some constant α , such that the following holds.*

1. *For all $a, a' \in P$, it holds that $\|a - a'\|_p > ((\zeta + 1)\gamma n)^{1/p}$.*
2. *For all $b, b' \in Q$, it holds that $\|b - b'\|_p > ((\zeta + 1)\gamma n)^{1/p}$.*
3. *For all $a \in P, b \in Q$, it holds that $\|a - b\|_p = (\zeta\gamma n)^{1/p}$.*

Moreover, there exists a deterministic algorithm that outputs P and Q in time $O(2^{n/2} n^{O(1)})$.

Proof. In literature, we note that for any constant $\delta > 0$, there is an explicit binary code of (some) constant rate and relative distance at least $\frac{1}{2} - \delta$ and the entire code can be listed in quasilinear time with respect to the size of the code (see Appendix E.1.2.5 from [Gol08], or Justesen codes [Jus72]). Using such a code, for some constant α , we obtain $C \subset \{0, 1\}^{\alpha n}$ of cardinality $2^{n/2}$, such that for every two $x, y \in C$, x and y differ on at least $(\frac{1}{2} - \delta)\alpha n$ coordinates, for some constant $\delta \in (0, \frac{1}{4} - \frac{1}{2^p})$. Moreover, C can be computed in $2^{n/2}n^{O(1)}$ time. From C , we obtain $C' \subset \{-h, h\}^{\alpha n}$ by replacing 0 with $-h$, and 1 with h in each of the coordinates in all the points of C . We construct the requisite P and Q as subsets of $\{-h, 0, h\}^{2\alpha n}$. For every $i \in [2^{n/2}]$, the i^{th} point of P is given by the concatenation of the i^{th} point of C' with $0^{\alpha n}$. Similarly, the i^{th} point of Q is given by the concatenation of $0^{\alpha n}$ with the i^{th} point of C' (note the reversal in the order of the concatenation). In particular, points in P and Q are of the form $(x_i, \vec{0})$ and $(\vec{0}, x_i)$, respectively, where x_i is the i^{th} point in C' and $\vec{0}$ is the zero-vector of length αn .

First, consider any two points in the same set, say $a, a' \in P$ (respectively, $b, b' \in Q$). We have from the distance of C that on at least $(\frac{1}{2} - \delta)\alpha n$ coordinates the two points differ by $2h$, thus implying that their L_p -distance is at least

$$\begin{aligned} \left(\left(\frac{1}{2} - \delta \right) \alpha n 2^p h^p \right)^{1/p} &> \left(\left(\frac{1}{4} + \frac{1}{2^p} \right) \alpha n 2^p h^p \right)^{1/p} \\ &= \left(\left(1 + \frac{1}{2^{p-2}} \right) \left(\frac{\gamma n}{1 - 2^{2-p}} \right) \right)^{1/p} \\ &= \left(\left(\frac{2^{p-2} + 1}{2^{p-2} - 1} \right) \gamma n \right)^{1/p} \\ &= ((\zeta + 1)\gamma n)^{1/p}. \end{aligned}$$

This proves the first two items of the Lemma. Next we prove the third item. Consider any two points from different sets, say $a \in P$ and $b \in Q$. It is easy to see from the construction that a and b differ in every coordinate by exactly h . Thus, the L_p -distance between any two points from different set is exactly

$$(2\alpha n h^p)^{1/p} = \left(\frac{2\gamma n}{2^{p-2} - 1} \right)^{1/p} = (\zeta \gamma n)^{1/p}. \quad \square$$

7 Connection to CLOSEST PAIR

In this section, we prove the quadratic-time hardness for CLOSEST PAIR in the L^∞ -metric and SET CLOSEST PAIR. Our reductions are from OV which we rephrase it as follows: Given a pair of collection of vectors $U, W \subseteq \{0, 1\}^d$, the goal is to find a pair of vectors $u \in U$ and $w \in W$ such that $(u_i, w_i) \in \{(0, 0), (0, 1), (1, 0)\}$ for all $i \in [d]$. Throughout, we denote by n the total number of vectors in U and W .

We start by giving a base reduction, which will be used in the hardness of CLOSEST PAIR in the L^∞ -metric and SET CLOSEST PAIR.

7.1 Base Reduction

Let $U, W \subseteq \{0, 1\}^d$ be an instance of OV. We may assume that U and W have no duplicates. Otherwise, we may sort vectors in U (resp., W) in lexicographic order and then sequentially remove duplicates; this preprocessing takes $O(dn \log n)$ -time.

We construct a pair of sets $A, B \subseteq \mathbb{R}^d$ of BCP from U, W as follows. For each vector $u \in U$ (resp., $w \in W$), we create a point $a \in A$ (resp., $b \in B$) such that

$$a_j = \begin{cases} 0 & \text{if } u_j = 0, \\ 2 & \text{if } u_j = 1. \end{cases} \quad b_j = \begin{cases} 1 & \text{if } w_j = 0, \\ -1 & \text{if } w_j = 1. \end{cases}$$

Observe that, for any vectors $a \in A$ and $b \in B$, $|a_j - b_j| = 2$ only if $a_j = b_j = 1$; otherwise, $|a_j - b_j| = 1$. It can be seen that $\|a - b\|_p = d$ if and only if their corresponding vectors $u \in U$ and $w \in W$ are orthogonal. Thus, this gives an alternate proof for the quadratic-time hardness of BCP under OVH.

7.2 Quadratic-Time Hardness of CLOSEST PAIR in L^∞

Here we show that the reduction in Section 7.1 rules out both exact and 2-approximation algorithm for CLOSEST PAIR in L^∞ that runs in subquadratic-time. (unless OVH is false). That is, we prove Theorem 4, which follows from the theorem below.

Theorem 25. *Assuming OVH, for any $\varepsilon > 0$ and $d = \omega(\log n)$, there is no $O(n^{2-\varepsilon})$ -time algorithm that, given a point-set $P \subseteq \mathbb{R}^d$, distinguishes between the following two cases:*

- *There exists a pair of vectors in P with L^∞ -distance one.*
- *Every pair of vectors in P has L^∞ -distance two.*

In particular, it is OVH-hard to approximate CLOSEST PAIR in the L^∞ -metric to within a factor of two.

Proof. Consider the sets A and B constructed from an instance of the OV problem in Section 7.1.

First, observe that every inner pair has L^∞ -distance at least 2. To see this, consider an inner pair $a, a' \in A$. Since all inner pairs are distinct, they must have at least one different coordinate, say $a_j \neq a'_j$ for some $j \in \{1, \dots, n/2\}$. Consequently, $(a_j, a'_j) \in \{(0, 2), (2, 0)\}$, implying that the L^∞ -distance of a and a' is at least 2. The case of an inner pair $b, b' \in B$ is similar. Thus, any pair of vectors with L^∞ -distance less than two must be a crossing pair $a \in A, b \in B$.

Now suppose there is a pair of orthogonal vectors $u^* \in U, w^* \in W$, and let $a^* \in A$ and $b^* \in B$ be the corresponding vectors of u^* and w^* in the CLOSEST PAIR instance, respectively. Then we know from the construction that $(a_j^*, b_j^*) \in \{(0, 1), (0, -1), (2, 1)\}$ for all coordinates $j \in [n]$. Thus, the L^∞ -distance of a^* and b^* must be one.

Next suppose there is no orthogonal pair of vectors in (U, W) . Then every vectors $u \in U$ and $w \in W$ must have one coordinate, say j , such that $u_j = w_j = 1$. So, the corresponding vectors a

and b (of u and w , respectively) must have $a_i = 2, b_j = -1$. This means that a and b have L^∞ -distance at least three. (Note that there might be an inner pair with L^∞ -distance two.) Therefore, we conclude that every pair of points in $A \cup B$ has L^∞ -distance at least two. \square

7.3 Quadratic-Time Hardness of SET CLOSEST PAIR

In this section, we consider SET CLOSEST PAIR in the L^p -metric with $p > 1$. The naive algorithm for this problem tries all possible pairs of subsets $a, b \in X$ and returns the smallest one. The running time of this trivial algorithm is thus $O(dtn^2)$. There is no better algorithm in literature, and we show that it might not be possible to improve the bound unless OVH is false.

More precisely, we rule out subquadratic-time algorithm for SET CLOSEST PAIR by a reduction from OVH, i.e., we prove Theorem 5. For an overview, we take a hardness instance (A, B) of BCP constructed in the base reduction in Section 7.1. Then we attach a codeword to each vector (possibly with duplicates) in A and B . This can be seen as a set *set version* of polar-pair.

Our construction requires the following functions.

1. Let $f_1 : \{0, 1\} \rightarrow \mathbb{R}^2$ be defined as $f_1(0) = (1, 0)$ and $f_1(1) = (-1, 0)$.
2. Let $f_{2,1} : \{0, 1\} \rightarrow \mathbb{R}^2$ be defined as $f_{2,1}(0) = (1, 1)$ and $f_{2,1}(1) = (1, -1)$.
3. Let $f_{2,2} : \{0, 1\} \rightarrow \mathbb{R}^2$ be defined as $f_{2,2}(0) = (-1, -1)$ and $f_{2,2}(1) = (-1, 1)$.

The functions $f_1, f_{2,1}$ and $f_{2,2}$ are illustrated in Figure 1.




Figure 1: An illustration of $f_1, f_{2,1}$ and $f_{2,2}$. The set $\{f_{2,1}(0), f_{2,2}(0)\}$ is colored by red and the set $\{f_{2,1}(1), f_{2,2}(1)\}$ is colored by blue.

We will abuse notation and define $f_* : \{0, 1\}^l \rightarrow \mathbb{R}^{2l}$ to be the function that maps a vector \vec{v} with l coordinates to a vector with $2l$ coordinates by applying f_* on each coordinate of \vec{v} .

7.3.1 Constructing Set-Version of Polar-Pair.

Consider the code from the proof of Lemma 24 (namely for $0 < \delta < \frac{1}{2} - \frac{1}{2^p}$. For some constant $\alpha > 0$, there is a binary code with constant rate and relative distance at least $\frac{1}{2} - \delta$. The code can be explicitly constructed in quasilinear-time with respect to the size of the code). Let $q = O(\log n)$ be the dimension of the mentioned code. We obtain a collection of codewords $C \subseteq \{0, 1\}^{\alpha q}$ of cardinality $2 \cdot 2^q$ such that, for every two $x, y \in C$, x and y differ on at least $(\frac{1}{2} - \delta) \alpha q$ coordinates.

Next we construct the sets \mathcal{P} and \mathcal{Q} so that

$$\mathcal{P} = \{\{f_1(c)\} : c \in C\} \quad \text{and} \quad \mathcal{Q} = \{\{f_{2,1}(c), f_{2,2}(c)\} : c \in C\}.$$

We use the following lemma which is a variant of Theorem 24.

Lemma 26. *The following properties hold for \mathcal{Q}, \mathcal{P} .*

1. *For all $X, X' \in \mathcal{P}$ it holds that $\text{dist}_p(X, X') > (2^p(1/2 - \delta))^{1/p} (\alpha q)^{1/p}$.*
2. *For all $Y, Y' \in \mathcal{Q}$ it holds that $\text{dist}_p(Y, Y') > (2^p(1/2 - \delta))^{1/p} (\alpha q)^{1/p}$.*
3. *For all $X \in \mathcal{P}, Y \in \mathcal{Q}$ it holds that $\text{dist}_p(X, Y) = (\alpha q)^{1/p}$.*

Moreover, there exists a deterministic algorithm that outputs \mathcal{P} and \mathcal{Q} in time $O(2^q n^{O(1)})$.

Proof. Let c and c' be any pair of codewords whose Hamming distance is $\eta \geq (1/2 - \delta) \alpha q$.

First, consider a pair of sets $X, X' \in \mathcal{A}$ which are attached with the codewords c and c' , respectively. Then we have

$$\begin{aligned} \text{dist}_p(X, X') &= \text{dist}_p(\{f_1(c_i)\}, \{f_1(c_j)\}) \\ &= \|f_1(c) - f_1(c')\|_p \\ &= (2^p \eta)^{1/p} \\ &> (2^p(1/2 - \delta))^{1/p} (\alpha q)^{1/p}. \end{aligned}$$

Next consider a pair of sets $Y, Y' \in \mathcal{B}$ which are attached with the codewords c and c' , respectively. Then we have

$$\begin{aligned} \text{dist}_p(Y, Y') &= (\eta \text{dist}_p(\{f_{2,1}(0), f_{2,2}(0)\}, \{f_{2,1}(1), f_{2,2}(1)\}))^{1/p} \\ &= (\eta 2^p)^{1/p} > (2^p(1/2 - \delta))^{1/p} (\alpha q)^{1/p}. \end{aligned}$$

Here the last equality follows by the definition of $d, f_1, f_{2,1}, f_{2,2}$.

Lastly, consider any crossing pair $X \in \mathcal{A}$ and $Y \in \mathcal{B}$. By definition, for every $k_1, k_2 \in \{0, 1\}$,

$$\text{dist}_p(\{f_1(k_1)\}, \{f_{2,1}(k_2), f_{2,2}(k_2)\}) = 1.$$

Hence, $\text{dist}_p(X, Y) = (\alpha q)^{1/p}$. □

7.3.2 Reduction to SET CLOSEST PAIR via Set-Polar-Pair

We take a pair of sets (A, B) from Section 7.1 and the code $(\mathcal{P}, \mathcal{Q})$ from Section 7.3.1. We assume wlog that $|A| = |\mathcal{P}|$, (resp., $|B| = |\mathcal{Q}|$) and match them arbitrary. For notational convenience, we denote by $\sigma(a)$ (resp., $\sigma(b)$) a set $X \in \mathcal{P}$ that is matched with a (resp., b).

Choose δ to be sufficiently small so that $2^p(1/2 - \delta) > 1$, and let M to be sufficiently large number, say $M \gg d^2$. By $z\sigma(a)$ (resp., $z\sigma(b)$), we mean a point-set obtained by scaling all the vectors in the set $\sigma(a)$ (resp., $\sigma(b)$) by a factor z . We construct a pair of collections of point-sets $(\mathcal{A}, \mathcal{B})$ which form an instance of SET CLOSEST PAIR as below.

$$\begin{aligned}\mathcal{A} &:= \left\{ a \times \frac{M}{\alpha q} \sigma(a) : a \in A \right\} \\ \mathcal{B} &:= \left\{ b \times \frac{M}{\alpha q} \sigma(b) : b \in B \right\}\end{aligned}$$

In other words, we attached a codeword to each vector $a \in A$ (resp., $b \in B$). The next claim will finish the proof of Theorem 5 (by setting $p > 1$).

Claim 27. *There exist two sets X, Y in $\mathcal{A} \cup \mathcal{B}$ such that $\text{dist}_p(X, Y)$ is at most $(M + d)^{1/p}$ if and only if there is an orthogonal pair $u \in U$ and $w \in W$.*

Proof. Observe that for any vectors $a \in A$ and $b \in B$, $|a_j - b_j| = 2$ only if $a_j = b_j = 1$; otherwise, $|a_j - b_j| = 1$.

First, suppose there is an orthogonal pair $u \in U$ and $w \in W$. Then there is a pair of sets $X \in \mathcal{A}$ and $Y \in \mathcal{B}$ corresponding to points $a \in A$ and $b \in B$ (in the base construction) such that for every $j \in [d]$, $|a_j - b_j|^p = 1$. This means that

$$\sum_{j=1}^d |a_j - b_j|^p = d.$$

By Lemma 26 and the construction of \mathcal{A} and \mathcal{B} (from \mathcal{P} and \mathcal{Q} respectively), we know that

$$\left(\frac{M}{\alpha q} \cdot \text{dist}_p(\{f_1(\sigma(a)), \{f_{2,1}(\sigma(b)), f_{2,2}(\sigma(b))\}\}) \right)^p = M.$$

Thus, we have

$$\text{dist}_p(X, Y) = \left(\sum_{j=1}^d |a_j - b_j|^p + \frac{M}{\alpha q} \cdot (\text{dist}_p(\{f_1(\sigma(a)), \{f_{2,1}(\sigma(b)), f_{2,2}(\sigma(b))\}\})^p \right)^{1/p} = (M + d)^{1/p}$$

Next suppose there is no orthogonal pair of vectors, i.e., for any pair of vectors $u \in U$ and $w \in W$, there is a coordinate j such that $u_j = w_j = 1$. Then any pair of vectors $a \in A$ and $b \in B$ must have one coordinate j such that $|a_j - b_j| = 2$. (Note that $|a_j - b_j| \in \{1, 2\}$ for all $j \in [d]$.)

Now consider a pair of sets $X, Y \in \mathcal{A} \cup \mathcal{B}$. We have the following cases.

- **Case 1:** $X, Y \in \mathcal{A}$. Let $a, a' \in A$ be vectors corresponding to X and Y , respectively. Then by the construction of \mathcal{A} and Lemma 26, we have that

$$\begin{aligned} \text{dist}_p(X, Y) &= \left(\sum_{j=1}^d |a_j - a'_j|^p + \left(\frac{M}{\alpha q} \text{dist}_p(\{f_1(\sigma(a)), \{f_{2,1}(\sigma(a')), f_{2,2}(\sigma(a'))\}\} \right)^p \right)^{1/p} \\ &> (2^p(1/2 - \delta)M^p)^{1/p} \\ &> (M + d)^{1/p} \end{aligned}$$

The last inequality follows because $M \gg d^2$ and $2^p(1/2 - \delta) > 1$.

- **Case 2:** $X, Y \in \mathcal{B}$. This case is exactly the same as the previous case, and we have that $\text{dist}_p(X, Y) > (M + d)^{1/p}$.
- **Case 3:** $X \in \mathcal{A}, Y \in \mathcal{B}$. Let $a \in A$ and $b \in B$ be points corresponding to X and Y , respectively. Then by the construction of X and Y , we have that

$$\begin{aligned} \text{dist}_p(X, Y) &= \left(\sum_{j=1}^d |a_j - b_j|^p + \frac{M}{\alpha q} \cdot (\text{dist}_p(\{f_1(\sigma(a)), \{f_{2,1}(\sigma(b)), f_{2,2}(\sigma(b))\}\})^p \right)^{1/p} \\ &\geq ((d - 1) + 2^p + M)^{1/p} \\ &> (M + d)^{1/p} \end{aligned}$$

Thus, in any cases, the L^p -distance between X and Y is strictly greater than $(M + d)^{1/p}$.

This concludes the claim. \square

8 Conclusion and Discussion

We have studied the sphericity and contact dimension of the complete bipartite graph in various metrics. We have proved new bounds on these measures for some metrics and given alternate proofs of the old bounds for the Euclidean metric. However, biclique sphericity and biclique contact dimension in the L^1 -metric remains poorly understood as we are unable to show any strong upper or lower bounds. However, we believe that both L^1 and L^2 metrics have linear upper and lower bounds. To be precise, we raise the following conjecture:

Conjecture 28 (L^1 -Biclique Sphericity Conjecture).

$$\text{bsph}(L^1) = \Omega(n).$$

In addition, we have showed conditional lower bounds for CLOSEST PAIR the L^p -metric, for all $p \in \mathbb{R}_{>2} \cup \{\infty\}$, by using polar-pairs of point-sets. However, it is unlikely that our techniques could be extended to popular metrics L^2 , L^1 , and L^0 . An open question is thus whether there exists an alternative technique to derive a lower bound from OVH to CLOSEST PAIR for these metrics. The answer might be on the positive side, i.e., there might exist an algorithm that performs well in the L^2 -metric because there are more tools available, e.g., Johnson-Lindenstrauss' dimension reduction. Thus, it is possible that there exists a strongly subquadratic-time algorithm in the L^2 -metric. This question is still mysterious and remains a long standing open problem.

Acknowledgements We would like to thank Petteri Kask and Rasmus Pagh for useful discussions and also for pointing out the reference [AW15]. We would like to thank Eylon Yogev and Amey Bhangale for some preliminary discussions. We would like to thank Uriel Feige for many discussions. Finally, we would like to thank Roei Tell for helping us improve the presentation of the paper.

9 References

- [AP03] Noga Alon and Pavel Pudlák. Equilateral sets in ℓ_p^n . *Geometric & Functional Analysis GAFA*, 13(3):467–482, 2003. 5
- [AW15] Josh Alman and Ryan Williams. Probabilistic polynomials and hamming nearest neighbors. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 136–150, 2015. 2, 4, 25
- [BCL98] Hans-Jürgen Bandelt, Victor Chepoi, and Monique Laurent. Embedding into rectilinear spaces. *Discrete & Computational Geometry*, 19(4):595–604, 1998. 5
- [Ben80] Jon Louis Bentley. Multidimensional divide-and-conquer. *Commun. ACM*, 23(4):214–229, 1980. 5
- [Ben83] Michael Ben-Or. Lower bounds for algebraic computation trees (preliminary report). In *Proceedings of the 15th Annual ACM Symposium on Theory of Computing, 25-27 April, 1983, Boston, Massachusetts, USA*, pages 80–86, 1983. 5
- [BL05] Yonatan Bilu and Nathan Linial. Monotone maps, sphericity and bounded second eigenvalue. *J. Comb. Theory, Ser. B*, 95(2):283–299, 2005. 2, 3
- [BS76] Jon Louis Bentley and Michael Ian Shamos. Divide-and-conquer in multidimensional space. In *Proceedings of the 8th Annual ACM Symposium on Theory of Computing, May 3-5, 1976, Hershey, Pennsylvania, USA*, pages 220–230, 1976. 5
- [CIP09] Chris Calabro, Russell Impagliazzo, and Ramamohan Paturi. The complexity of satisfiability of small depth circuits. In *Parameterized and Exact Computation, 4th International Workshop, IWPEC 2009, Copenhagen, Denmark, September 10-11, 2009, Revised Selected Papers*, pages 75–85, 2009. 4
- [DM94] Michel Deza and Hiroshi Maehara. A few applications of negative-type inequalities. *Graphs and Combinatorics*, 10(2-4):255–262, 1994. 10
- [FM88] Peter Frankl and Hiroshi Maehara. On the contact dimensions of graphs. *Discrete & Computational Geometry*, 3:89–96, 1988. 2, 3
- [Gol08] Oded Goldreich. *Computational Complexity: A Conceptual Perspective*. Cambridge University Press, New York, NY, USA, 1 edition, 2008. 19
- [GS16] Omer Gold and Micha Sharir. Dominance products and faster algorithms for high-dimensional closest pair under ℓ_∞ . *CoRR*, abs/1605.08107, 2016. 5

- [Guy83] Richard K Guy. An olla-podrida of open problems, often oddly posed. *The American Mathematical Monthly*, 90(3):196–200, 1983. [5](#)
- [HNS88] Klaus H. Hinrichs, Jürg Nievergelt, and Peter Schorn. Plane-sweep solves the closest pair problem elegantly. *Inf. Process. Lett.*, 26(5):255–261, 1988. [5](#)
- [ILLP04] Piotr Indyk, Moshe Lewenstein, Ohad Lipsky, and Ely Porat. Closest pair problems in very high dimensions. In *Automata, Languages and Programming: 31st International Colloquium, ICALP 2004, Turku, Finland, July 12-16, 2004. Proceedings*, pages 782–792, 2004. [5](#)
- [IP01] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001. Preliminary version in CCC’99. [4](#)
- [IPZ01] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001. Preliminary version in FOCS’98. [4](#)
- [Jus72] Jørn Justesen. Class of constructive asymptotically good algebraic codes. *IEEE Trans. Information Theory*, 18(5):652–656, 1972. [19](#)
- [KLS00] Jack H. Koolen, Monique Laurent, and Alexander Schrijver. Equilateral dimension of the rectilinear space. *Des. Codes Cryptography*, 21(1/3):149–164, 2000. [5](#)
- [KM95] Samir Khuller and Yossi Matias. A simple randomized sieve algorithm for the closest-pair problem. *Inf. Comput.*, 118(1):34–37, 1995. [5](#)
- [Mae84] Hiroshi Maehara. Space graphs and sphericity. *Discrete Applied Mathematics*, 7(1):55–64, 1984. [2](#)
- [Mae85] Hiroshi Maehara. Contact patterns of equal nonoverlapping spheres. *Graphs and Combinatorics*, 1(1):271–282, 1985. [2](#), [3](#)
- [Mae91] Hiroshi Maehara. Dispersed points and geometric embedding of complete bipartite graphs. *Discrete & Computational Geometry*, 6:57–67, 1991. [2](#), [3](#)
- [Rab76] Michael O. Rabin. Probabilistic algorithms. In *Proceedings of a Symposium on New Directions and Recent Results in Algorithms and Complexity, Computer Science Department, Carnegie-Mellon University, April 7-9, 1976*, pages 21–39, 1976. [5](#)
- [Rob69] Fred S Roberts. On the boxicity and cubicity of a graph. *Recent Progresses in Combinatorics*, pages 301–310, 1969. [2](#)
- [SH75] Michael Ian Shamos and Dan Hoey. Closest-point problems. In *16th Annual Symposium on Foundations of Computer Science, Berkeley, California, USA, October 13-15, 1975*, pages 151–162, 1975. [5](#)
- [Ver10] Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. *CoRR*, abs/1011.3027, 2010. [17](#)

- [Wil05] Ryan Williams. A new algorithm for optimal 2-constraint satisfaction and its implications. *Theor. Comput. Sci.*, 348(2-3):357–365, 2005. Preliminary version in ICALP’04. [4](#), [6](#)
- [Wil15] Virginia Vassilevska Williams. Hardness of easy problems: Basing hardness on popular conjectures such as the strong exponential time hypothesis (invited talk). In *10th International Symposium on Parameterized and Exact Computation, IPEC 2015, September 16-18, 2015, Patras, Greece*, pages 17–29, 2015. [6](#)
- [Wil16] Virginia Vassilevska Williams. Fine-grained algorithms and complexity (invited talk). In *33rd Symposium on Theoretical Aspects of Computer Science, STACS 2016, February 17-20, 2016, Orléans, France*, pages 3:1–3:1, 2016. [6](#)
- [Yao91] Andrew Chi-Chih Yao. Lower bounds for algebraic computation trees with integer inputs. *SIAM J. Comput.*, 20(4):655–668, 1991. Preliminary version in FOCS’89. [5](#)